

1 1. (Amended) A method for processing a database query, comprising:
2 partially pre-aggregating records in a database according to a single
3 grouping column to provide a result that contains at least two records having like
4 grouping column values; and
5 aggregating records derived from the partial pre-aggregation to provide a
6 result that contains records having unique grouping column values.

a'
7
8 2. The method as recited in claim 1, wherein the partially pre-
9 aggregating further comprises:
10 maintaining a record store in memory, the record store having one record
11 for each different grouping column value encountered in the operation;
12 receiving a new record;
13 combining the new record with a record having the same grouping column
14 value, if such a record exists; and
15 adding the new record to the record store in the memory if there is no record
16 in the record store that has the same grouping column value as the new record.

17
18 3. The method as recited in claim 2, further comprising:
19 adding additional new records to the record store until the record store
20 reaches a capacity such that it can accept no new records; and
21 outputting one or more records from the record store to a subsequent
22 database operator.
23
24
25

a' 1 4. The method as recited in claim 3, wherein after the one or more
2 records have been output to the subsequent database operator, the adding and
3 outputting are repeated until there are no new records to process.

4
5 5. The method as recited in claim 4, wherein any records remaining in
6 the record store after there are no new records to process are output to the
7 subsequent database operator.

8
9 6. The method as recited in claim 3, wherein the subsequent database
10 operator is a join.

11
12 7. The method as recited in claim 1, further comprising estimating the
13 costs and benefits of the partial pre-aggregation, and partially pre-aggregating the
14 records only if the estimating indicates that the benefits are greater than the costs.

15
16 8. The method as recited in claim 1, wherein the partially pre-
17 aggregating includes utilizing a hashing function.

18
19 9. The method as recited in claim 1, wherein the partial pre-aggregating
20 creates a record store in memory, and wherein the method further comprises
21 utilizing the record store in memory for one or more other database operators.

22
23 10. One or more computer-readable media having computer-executable
24 instructions that, when executed by a computer, perform the method recited in
25 claim 1.

1
2 11. A computer programmed to perform the method recited in claim 1.

3
4 12. (Amended) A relational database system, comprising:
5 memory for storing a record store, the memory having a portion available
6 for query processing;

a' 7 a query processor coupled to the memory to process a query on the record
8 store according to a single grouping column, the query processor being configured
9 to partially pre-aggregate the record store to provide a result that contains at least
10 two data records that have like grouping column values; and

11 the query processor being further configured to aggregate data records
12 resulting from the partial pre-aggregation to provide an aggregation result that
13 contains data records, no two of the data records having the same grouping column
14 value.

15
16 13. (Amended) The relational database system as recited in claim 12,
17 wherein the query processor being configured to partially pre-aggregate the record
18 store further comprises the query processor being configured to:

19 maintain a record store in the volatile memory, the record store having one
20 record for each different grouping column value encountered in the partial pre-
21 aggregation;

22 receive an input record from the non-volatile memory;

23 combine the input record with a record in the record store that has the same
24 grouping column value, if there is such a record; and
25

1 adding the input record to the record store if there is no record in the record
2 store that has the same grouping column value as the input record.

3
4 14. The relational database system as recited in claim 13, wherein the
5 query processor is further configured to:

6 add additional input records to the record store in memory until the record
7 store reaches a capacity such that it can accept no more input records; and
8 output the records in the record store to a subsequent database operator.

9
10 15. The relational database system as recited in claim 14, wherein the
11 query processor is configured to:

12 continually add input records to the record store; and
13 output one or more records from the record store to a subsequent database
14 operator when the record store reaches a capacity such that it can accept no more
15 new records, or whenever there are no new records to process.

16
17 16. The relational database system as recited in claim 12, wherein the
18 query processor is further configured to perform a join on the records resulting
19 from the partial pre-aggregation.

20
21 17. The relational database system as recited in claim 12, wherein the
22 query processor is further configured to create a record store as a result of the
23 partial pre-aggregation and utilize the record store in processing of another
24 database operator.

1 18. The relational database system as recited in claim 12, further
2 comprising:

3 a query optimizer configured to estimate the costs and benefits of the query
4 processor performing a partial pre-aggregation; and

5 wherein the query processor performs the partial pre-aggregation only if the
6 query optimizer indicates that the benefits of the partial pre-aggregation are greater
7 than the costs of the partial pre-aggregation.

8
9 19. The relational database system as recited in claim 12, wherein the
10 query processor is further configured to utilize hashing to perform the partial pre-
11 aggregation.

12
13 20. The relational database system as recited in claim 12, wherein the
14 query processor is further configured to utilize hashing and partitioning to perform
15 the partial pre-aggregation.

16
17 21. (Amended) A relational database computer program embodied on a
18 computer-readable medium, comprising:

19 partial pre-aggregation code to partially pre-aggregate data records
20 according to grouping column values in a single grouping column to provide a
21 partial pre-aggregation result having two or more records having like grouping
22 column values; and

23 aggregation code to aggregate data records in the partial pre-aggregation
24 result to provide an aggregation result having records with unique grouping
25 column values.

1
2 22. The relational database computer program as recited in claim 21,
3 wherein the partial pre-aggregation code is designed to:

4 maintain a record store in memory, the record store having one record for
5 each different grouping column value encountered in the partial pre-aggregation;

6 receive a new record;

7 combine the new record with a record in the record store having the same
8 grouping column value, if such a record exists;

9 add the new record to the record store no record in the record store that has
10 the same grouping column value as the new record;

11 continuously add additional new records to the record store until the record
12 store has reached a record store capacity;

13 output one or more records from the record store to a subsequent database
14 operator when the record store has reached the record store capacity; and

15 output the records in the record store when there are no new records to
16 process.

17
18 23. The relational database computer program as recited in claim 22,
19 further comprising database operator code that utilizes the record store for input.
20
21
22
23
24
25

24. (Amended) A relational database computer program stored on a computer-readable medium, the relational database computer program comprising computer-executable instructions that, when executed on a computer, perform the following steps:

receiving a stream of input records;

aggregating the input records in the stream according to a single grouping column as it is received to create a record store;

joining records in the record store with other data;

outputting the records in the record store after the join;

aggregating the records output from the join; and

wherein the records output from the join include at least two records that have an identical grouping column value in the single grouping column.

25. The relational database computer program as recited in claim 24, wherein:

the record store has a capacity that is less than the number of records in the stream of input records; and

the aggregating each input record is performed until the record store reaches capacity.

26. The relational database computer program as recited in claim 24, further comprising computer-executable instructions that, when executed by a computer, perform the following steps:

determining if it is optimal to aggregate the input records prior to performing the join; and

1 performing the aggregation prior to the join only if a determination is made
2 that it is optimal to perform an aggregation prior to the join.

3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25